

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

5

APPLICATION PAPERS

10

OF

FRASER PETER HOWARD

15

ANDREW KEMP

ALEX JAMES HINCHLIFFE

20

AND

BOBBY RAI

25

FOR

CENTRALLY MANAGED MALWARE SCANNING

30

BACKGROUND OF THE INVENTION

Field of the Invention

5

This invention relates to the field of data processing systems. More particularly, this invention relates to the field of malware scanning, such as, for example, scanning for computer viruses, Trojans, worms, banned computer files or E-mails containing banned words or content.

10

Description of the Prior Art

It is known to provide malware scanning systems and mechanisms for identifying malware within a computer file to be accessed. As the number of computer viruses and the like that are present in the wild increases the processing associated with scanning a computer file to identify the presence of any of those viruses similarly increases. This increase in the processing required is disadvantageous. Furthermore, with the increasing levels of network connectivity between computer systems and the use of E-mail and other fast messaging systems, the spread of computer viruses has become increasingly rapid. In a number of recent cases mass mailing viruses have spread at such a speed that considerable damage has been caused before appropriate countermeasures have been able to be put in place. The delay in deploying such countermeasures is further increased by the need in many cases to update virus definition data on every individual node computer to be protected. The speed of spread of recent viruses is such that the download and installation delays associated with installing such new countermeasures, even when they have been developed and are available, is a significant disadvantage.

An anti-virus system produced by Sophos uses a locally held record of previously conducted on-access scans on the individual computer in question to determine whether or not a scan should be conducted for a computer file or a previously determined result used instead when this is available.

SUMMARY OF THE INVENTION

Viewed from one aspect the present invention provides a computer program product for controlling a computer to detect malware, said computer program product comprising:

detecting logic operable to detect a file access request to a computer file by a requesting computer;

10 file access clearance request generating logic operable to generate a file access clearance request including data identifying said computer file;

file access clearance request transmitting logic operable to transmit said file access clearance request from said requesting computer to an assessment computer responsible for assessment of whether said computer file contains malware;

15 file access clearance response receiving logic operable to receive at said requesting computer a file access clearance response from said assessment computer; and

file access permitting logic operable if said file access clearance response indicates said computer file does not contain malware, to permit said file access request by said requesting computer.

20 The invention recognises that an individual computer requiring a file access may pass the task of determining whether or not that access should be allowed to a different assessment computer. Whilst at first sight this may seem that it would slow down the file access, in practice there are considerable advantages. For example, the
25 assessment computer to which the task is passed may be more rapidly updated with new virus definition data as this is released than would be possible for the requesting computer. Accordingly, the requesting computer can benefit from the most up-to-date virus definition data more rapidly. Furthermore, the invention is particularly well suited to systems in which a plurality of requesting computers share an assessment computer since in many cases the individual requesting computers will show a high degree of correlation in the computer files to which they are requesting access and for which a malware scan is needed. Thus, rather than each requesting computer

individually scanning the computer files that are also being scanned by a large number of other computers, (e.g. the computer files associated with the operating system shared by all the requesting computers), these computer files can instead be scanned once by the assessment computer and then each requesting computer can check that the computer file has been scanned and cleared by sending an appropriate request to the assessment computer rather than having to scan the file itself. Thus, in exchange for the use of a small amount of bandwidth on the connections between the requesting computers and the assessment computer, a considerable processing burden may be lifted from the requesting computers. The additional processing burden on the assessment computer does not increase disproportionately since the high degree of correlation between the computer files accessed by the different requesting computers means that in many cases the assessment computer will be able to respond to a clearance request from a requesting computer on the basis that it has already scanned that file and without the need to rescan that file.

15

In order to identify the computer file to be scanned in a secure manner to the assessment computer, the data identifying the computer file preferably includes a checksum value. Checksum values may be made sufficiently specific for a computer file from which they have been calculated so as to be difficult to bypass as a way of uniquely identifying a computer file.

Additional data that is highly useful to the assessment computer in managing the requests it receives includes the filename of the computer file, data identifying the requesting computer and the storage location of the computer file. This provides useful audit information as well as providing for the possibility of files having the same name and same storage location on different computers in fact being different files with different checksums, both of which should be treated as separate entities when determining whether or not a request from another requesting computer relates to a file that has already scanned.

25
30

In circumstances when a requesting computer wishes to access a file that has not already been scanned, the assessment computer may send a scan request message

back to the requesting computer such that the requesting computer may send a copy of the file to the assessment computer for scanning. It will be appreciated that this may be slower than scanning the file locally, but the benefit will be that should a different requesting computer later wish to access that same file, then the scanning result produced by this action can be shared with that other computer, so speeding the operation of that other computer.

If the assessment computer determines that access should be denied to a computer file, then this may be used to trigger denied access actions within the requesting computer, the assessment computer, or elsewhere. Such denied access actions may include deletion of the computer file, repair of the computer file, quarantining of the computer file, generation of user warning messages, generation of administrator warning messages and the like.

Viewed from another aspect the invention also provides a computer program product for controlling a computer to detect malware, said computer program product comprising:

file access request receiving logic operable to receive at an assessment computer a file access clearance request from a requesting computer, said file access clearance request including data identifying a computer file to be accessed by said requesting computer;

file access clearance response generating logic operable in dependence upon said data identifying said computer file to determine if said computer file has previously been assessed as not containing malware and to generate a file access clearance response; and

file access clearance response transmitting logic operable to transmit said file access clearance response to said requesting computer.

The assessment computer may advantageously store the results of its previous scans within a database. This database may include a variety of fields relating to each computer file that has been scanned. These fields may include an access flag indicating whether access is to be denied to that computer file and a persistence flag

indicating whether or not the entry for that computer file should be purged during purge operations.

The access flag is particularly useful as in addition to allowing the recording
5 that access should be denied due to malware being detected, it also allows the central
management of whether or not a particular individual file or class or type of file should
be permitted to be accessed by all of those requesting computers that seek their access
permissions from that assessment computer. This centralised control is a powerful tool
that may be used to implement techniques such as the triggering of a lock down mode
10 of operation in which higher level security provisions are put in place by denying
access to certain files or types of files. As an example, if a message was received
indicating that the higher security mode should be entered, then the assessment
computer may use its mechanisms to deny access to any newly encountered computer
file that had not previously been scanned and cleared for use. This would typically
15 allow the large majority of computer activity to continue whilst providing protection
against newly released malware threats until the appropriate countermeasures could be
put in place.

The persistence flag allows control of the flushing of entries from the
20 assessment computer. Whilst one of the advantages of the invention is storing the
results of previously conducted scans such that they need not be repeated, this has to
be tempered by allowing the results to be refreshed at a later time for at least some
files. It is possible to envisage that a particular computer file carrying a newly
25 released virus may not be detected as carrying that virus when it is first scanned, but
later when a new virus driver is available, that computer file would be detected and
blocked. Accordingly, as an example, it may be that all previous scan results could be
purged from the system whenever the virus definition data was updated.

Viewed from a further aspect the invention provides a computer program
30 product for controlling a computer to detect malware, said computer program product
comprising:

file access request detecting logic operable to detect a file access request to a computer file by a requesting computer;

file access clearance request generating logic operable to generate a file access clearance request including data identifying said computer file;

5 file access clearance request transmitting logic operable to transmit said file access clearance request from said requesting computer to an assessment computer responsible for assessment of whether said computer file contains malware;

file access clearance request receiving logic operable to receive at said assessment computer said file access clearance request from a requesting computer;

10 file access clearance response generating logic operable in dependence upon said data identifying said computer file to determine if said computer file has previously been assessed as not containing malware and to generate a file access clearance response;

file access clearance response transmitting logic operable to transmit said file 15 access clearance response to said requesting computer;

file access clearance response receiving logic operable to receive at said requesting computer said file access clearance response from said assessment computer; and

20 file access permitting logic operable if said file access clearance response indicates said computer file does not contain malware to permit said file access request by said requesting computer.

As well as the complementary aspects of the invention embodied in the form of the client software, the server software and the combination of the client and server 25 software, the invention may also take the form of corresponding methods of malware detection and apparatus for malware detection.

The above, and other objects, features and advantages of this invention will be apparent from the following detailed description of illustrative embodiments which is to 30 be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 schematically illustrates a computer network to which the present
5 technique may be applied;

Figure 2 is a flow diagram illustrating an example of the processing that may be
performed on a client computer;

10 Figure 3 is a flow diagram illustrating an example of the processing that may be
performed on a server computer;

Figure 4 is a flow diagram illustrating an example process that may be run on a
server computer waiting for a lock down signal;

15 Figure 5 schematically illustrates a database entry relating to a computer file that
has previously been scanned; and

20 Figure 6 is a diagram schematically illustrating the form of a general purpose
computer that may be used to implement the above described techniques.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 schematically illustrates a computer network 2 comprising a plurality of
25 client computers 4, 6, 8, 10, 12, 14, 16, 18 connected to respective departmental servers
20, 22 and to a corporate anti-virus (assessment) server 24. The computer network 2 is
also connected via a network link to an anti-virus providers FTP server 26 from which
virus definition data may be downloaded and from which messages, such as lock down
messages or messages relating to the availability of new virus definition data may be
30 received.

In operation a client computer 8 conducts processing activity during which it seeks to make a file access to a particular computer file (in this example, winword.exe). Rather than scanning this file itself for malware such as computer viruses, Trojans, worms, banned content etc, the client computer 8 instead passes an access clearance 5 request relating to the computer file through to the assessment computer 24 via the departmental server 20. The access clearance request includes the file name of the computer file being accessed, a checksum derived from that file in an effort to uniquely identify it (e.g. an MD5 checksum), data identifying the name of the client computer 8, and the path used by that client computer 8 to access the computer file. The assessment 10 computer 24 receives the access clearance request and uses the data identifying the computer file to look up within an associated database 28 as to whether or not a malware scan has already been conducted for that particular computer file. The computer file is uniquely identified by its filename and checksum value. If the computer file in question has already been scanned, then the result of that scan may be reused rather than 15 conducting the scan again. The assessment computer 24 can accordingly pass back to the client computer 8 a clearance request response indicating whether or not access to that file is permitted or some other further action should be taken. The client computer 8 can then use this access clearance response to either permit access to the file or take the further action specified.

20

The further action specified could be that the computer file has already been identified as containing malware and accordingly an appropriate anti-virus or other anti-malware response should be initiated on the client computer 8. A corresponding response could already have been initiated on the assessment computer 24. The denied access 25 responses could take the form of deleting the computer file concerned, repairing the computer file concerned, quarantining the computer file concerned, issuing a user alerting message, issuing an administrator alerting message, issuing a message to an anti-virus provider or some other action.

30

A different type of further action that could be required by the client computer 8 in response to the access clearance response arises when the assessment computer 24 has not previously scanned that file and accordingly the client computer 8 should send a copy

of that file through to the assessment computer 24 to be scanned. When such a copy of the computer file has been sent through to the assessment computer 24 and scanned, an appropriate pass or denied access action triggering response can be sent back from the assessment computer 24 to the client computer 8.

5

It will be appreciated that the client computers 4, 6, 8, 10, 12, 14, 16, 18 operate to share the results of the malware scanning that is performed on their behalf by the assessment computer 24. Since there is likely to be a high degree of correlation between the files being accessed by the different client computers, in many cases the scan that a client computer is requesting will already have been performed on behalf of another computer and that scan need not be run again by the assessment computer but instead merely the result of the previous scan returned. The use of a checksum allows the computer files to be reliably uniquely identified.

15

The co-ordination of file access permissions by the assessment computer 24 also allows central management of which computer files may be accessed upon the network 2 in a manner that allows the rapid implementation of any higher security level mode should this be desired. As an example, a lock down trigger message may be received from the anti-virus provider FTP server 26 that will trigger the assessment computer 24 to enter a higher security level mode compared to its normal mode of operation. In the higher security level mode, it may be that the access to whole variety of different types of files may be temporarily banned. As an example, the access to VBS files which are often a source of viruses may be banned across the network 2, as could access to e-mail attachments which are another potential source of virus propagation. Thus, when a new threat is identified by the anti-virus provider, lock down messages may be used to trigger predetermined, in accordance with user configuration and preferences, higher security level modes within connected assessment computers 24 to provide a degree of protection for the networks concerned whilst allowing the majority of their normal operations to continue.

25

30

Figure 2 is a flow diagram illustrating the processing performed by a client computer. At step 30 the computer waits to receive a file access request from a computer

program executing on that client computer. At step 32 a checksum value is calculated in accordance with one of several different possible checksum algorithms, such as the MD5 algorithm. At step 34 the filename, path, originating computer name and checksum value are sent through to the assessment computer. At step 36 the client computer waits to receive a response from the assessment computer. When that response is received, step 38 serves to determine whether or not the response indicates that the computer file had passed its scanning. If the computer file had passed its scanning and access is to be permitted, then processing proceeds to step 40 at which the access is allowed.

10 If the response is not that the computer file had passed its scan, then processing proceeds to step 42 at which a test is made as to whether or not the response indicated that a remote scan was required. If a remote scan was required, then the computer file in question is sent to the assessment computer at step 44 and processing is returned to step 36.

15 If the test at step 42 does not identify a remote scan request, then since the response is neither a pass, or a remote scan request, then the computer file must be one to which access is denied and accordingly step 46 serves to trigger the denied access actions. These may include deletion, repair, quarantining or other actions upon the computer file in question as well as the generation of appropriate warning messages to a user or an administrator.

20 Figure 3 illustrates the processing performed by the assessment computer. At step 48 the assessment computer waits to receive a request from a client computer. When a request is received, step 50 serves to identify whether a copy of a computer file is being returned to the assessment computer for scanning following a remote scan request that had earlier been issued by the assessment computer. If a copy of a computer file is being returned for scanning, then processing proceeds to step 52 at which the necessary malware scanning is conducted. Step 54 determines whether or not the computer file passed this malware scanning. If the computer file did pass the malware scanning, then step 56 serves to add details of that computer file to the database of scanned files held at

the assessment computer relating to scans already performed. Step 58 then returns the pass result to the client computer.

If the scanning at step 52 was detected as not being passed at step 54, then
5 processing proceeds to step 60 at which malware detected actions are triggered within the assessment computer. These malware detected actions may be similar to those previously described in relation to the client computer file. In addition, an entry specifying the malware scan fail of that computer file may also be added to the database. At step 62 an access denied result is returned to the client computer.

10

If the test at step 50 did not indicate that a computer file was being returned for scan, then step 64 serves to compare the file details being passed to the server with the database entries of computer files that have already been scanned. The computer file may be uniquely identified by its filename and its checksum. If at step 66 a match within the
15 database is detected, then step 68 determines whether or not this entry indicates that access should be allowed to that computer file. If access is not to be allowed, then processing proceeds to step 60. If access is allowed, then processing proceeds to step 70 at which a pass result is returned to the client computer.

20

If the test at step 66 indicated that a match was not found within the database, then step 72 serves to return a remote scan required result to the client computer in order to trigger the client computer to return a copy of the computer file to the assessment computer for scanning at step 52.

25

Figure 4 illustrates a process that may run on the assessment computer as a background task. At step 74 the assessment computer waits to receive a lock down trigger message from the anti-virus provider. If such a message is received, then step 76 serves to activate a lock down mode in the assessment computer. The lock down mode can switch on a user predetermined set of measures intended to provide a higher degree
30 of security, normally at the cost of at least some functionality. As an example, access to e-mail attachments or VBS files as a class may be denied. These are known to be particular vulnerabilities.

At step 78 the assessment computer waits for a user command to cancel the lock down mode. When such a command is received, then step 80 serves to cancel the lock down mode.

5

Figure 5 schematically illustrates the data that may be stored for a particular computer file within the database of previously conducted scan results held by the assessment computer 24. For an individual computer file, its filename, the originating computer for which the scan of that file was first conducted, the path to the file, the 10 checksum value for the file, an allowed access flag and a persistence flag are all stored. The allowed access flag may be used to indicate whether or not that file passed its scan result. The allowed access flag may also be used as a powerful tool for switching off or on access to individual files or classes of files by an administrator. The persistence flag controls how the entry is flushed from the database on a regular interval, such as when 15 new virus definition data is received. It will be appreciated that other fields could be added to the database relating to the particular file as required.

Figure 6 schematically illustrates a general purpose computer 200 of the type that may be used to implement the above techniques. The general purpose computer 200 20 includes a central processing unit 202, a random access memory 204, a read only memory 206, a hard disk drive 208, a display driver 210 and display 212, a user input/output circuit 214 and keyboard 216 and mouse 218 and a network interface unit 220 all connected via a common bus 222. In operation the central processing unit 202 executes program instructions stored within the random access memory 204, the read only 25 memory 206 or the hard disk drive 208. The working memory is provided by the random access memory 204. The program instructions could take a variety of forms depending on the precise nature of the computer 200 and the programming language being used. The results of the processing are displayed to a user upon the display 212 driven by the display driver 210. User inputs for controlling the general purpose 30 computer 200 are received from the keyboard 216 and the mouse 218 via the user input/output circuit 214. Communication with other computers, such as exchanging e-

mails, downloading files or providing internet or other network access, is achieved via the network interface unit 220.

It will be appreciated that the general purpose computer 200 operating under
5 control of a suitable computer program may perform the above described techniques and provide apparatus for performing the various tasks described. The general purpose computer 200 also executes the method described previously. The computer program product could take the form of a recordable medium bearing the computer program, such as a floppy disk, a compact disk or other recordable medium. Alternatively, the computer
10 program could be dynamically downloaded via the network interface unit 220.

It will be appreciated that the general purpose computer 200 is only one example of the type of computer architecture that may be employed to carry out the above described techniques. Alternative architectures are envisaged and are capable of use with
15 the above described techniques.

Although illustrative embodiments of the invention have been described in detail herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various changes and
20 modifications can be effected therein by one skilled in the art without departing from the scope and spirit of the invention as defined by the appended claims.